



US008489694B2

(12) **United States Patent**
Bhogal et al.

(10) **Patent No.:** **US 8,489,694 B2**
(45) **Date of Patent:** **Jul. 16, 2013**

(54) **PEER-TO-PEER COLLABORATION OF PUBLISHERS IN A PUBLISH-SUBSCRIPTION ENVIRONMENT**

(75) Inventors: **Kulvir S. Bhogal**, Pflugerville, TX (US); **Martin J. Gale**, Hampshire (GB); **Rick A. Hamilton, II**, Charlottesville, VA (US); **Andrew R. Jones**, Round Rock, TX (US); **Brian M. O'Connell**, RTP, NC (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 200 days.

(21) Appl. No.: **13/033,639**

(22) Filed: **Feb. 24, 2011**

(65) **Prior Publication Data**
US 2012/0221649 A1 Aug. 30, 2012

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.**
USPC **709/206**

(58) **Field of Classification Search**
USPC 709/206
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,406,537 B2	7/2008	Cullen	
7,574,408 B2	8/2009	Dominowska et al.	
7,698,372 B2	4/2010	Codignotto	
2003/0115291 A1	6/2003	Kendall et al.	
2005/0198109 A1*	9/2005	Teodosiu et al.	709/201
2006/0047666 A1	3/2006	Bedi et al.	
2006/0146999 A1	7/2006	Thompson et al.	
2006/0149840 A1	7/2006	Thompson et al.	

2006/0167897 A1	7/2006	Jones et al.	
2006/0168331 A1	7/2006	Thompson et al.	
2007/0061487 A1*	3/2007	Moore et al.	709/246
2007/0174232 A1	7/2007	Barcia et al.	
2008/0133541 A1	6/2008	Fletcher et al.	
2009/0177753 A1	7/2009	Banks et al.	
2009/0228563 A1	9/2009	Jones et al.	
2010/0162268 A1	6/2010	Thomas et al.	

FOREIGN PATENT DOCUMENTS

CN	101035093 A	9/2007
CN	101193077 A	6/2008
WO	2010/145362 A1	12/2010

OTHER PUBLICATIONS

Yan Li, PCT Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority, or the Declaration, International Application No. PCT/IB2012/050805, Date of Mailing Jul. 12, 2012, 4 pages.
Maitland, J., "Keeping Control Isn't Easy", Chapter 4: Cloud-Based Infrastructure, SearchCloudComputing.com, 13 pages.
Mell et al., "The NIST Definition of Cloud Computing", National Institute of Standards and Technology, Information Technology Laboratory, Version 15, Oct. 7, 2009, 2 pages.

* cited by examiner


Primary Examiner — Ryan Jakovac

(74) Attorney, Agent, or Firm — William E. Schiesser; Keohane & D'Alessandro, PLLC

(57) **ABSTRACT**

Embodiments of the present invention provide an approach for an enhanced pub/sub environment whereby a peer group of publishers is defined and can collaborate to determine when to cache data being sent to a common broker or brokers. Specifically, in a typical embodiment, configuration data or the like will be read to define a peer group among the publisher systems in the pub/sub messaging environment. Once such a group has been identified, communication therebetween will be established. Once cacheable data has been identified, a unique identifier/set of cache keys will be generated corresponding thereto and associated with the set of data. A message comprising the set of data and corresponding unique identifier will then be sent to a broker system (e.g., from the publisher system(s)), which will store (or have stored) the set of data in a cache, and pass the underlying message along to applicable subscriber system(s).

17 Claims, 6 Drawing Sheets



10 ↗

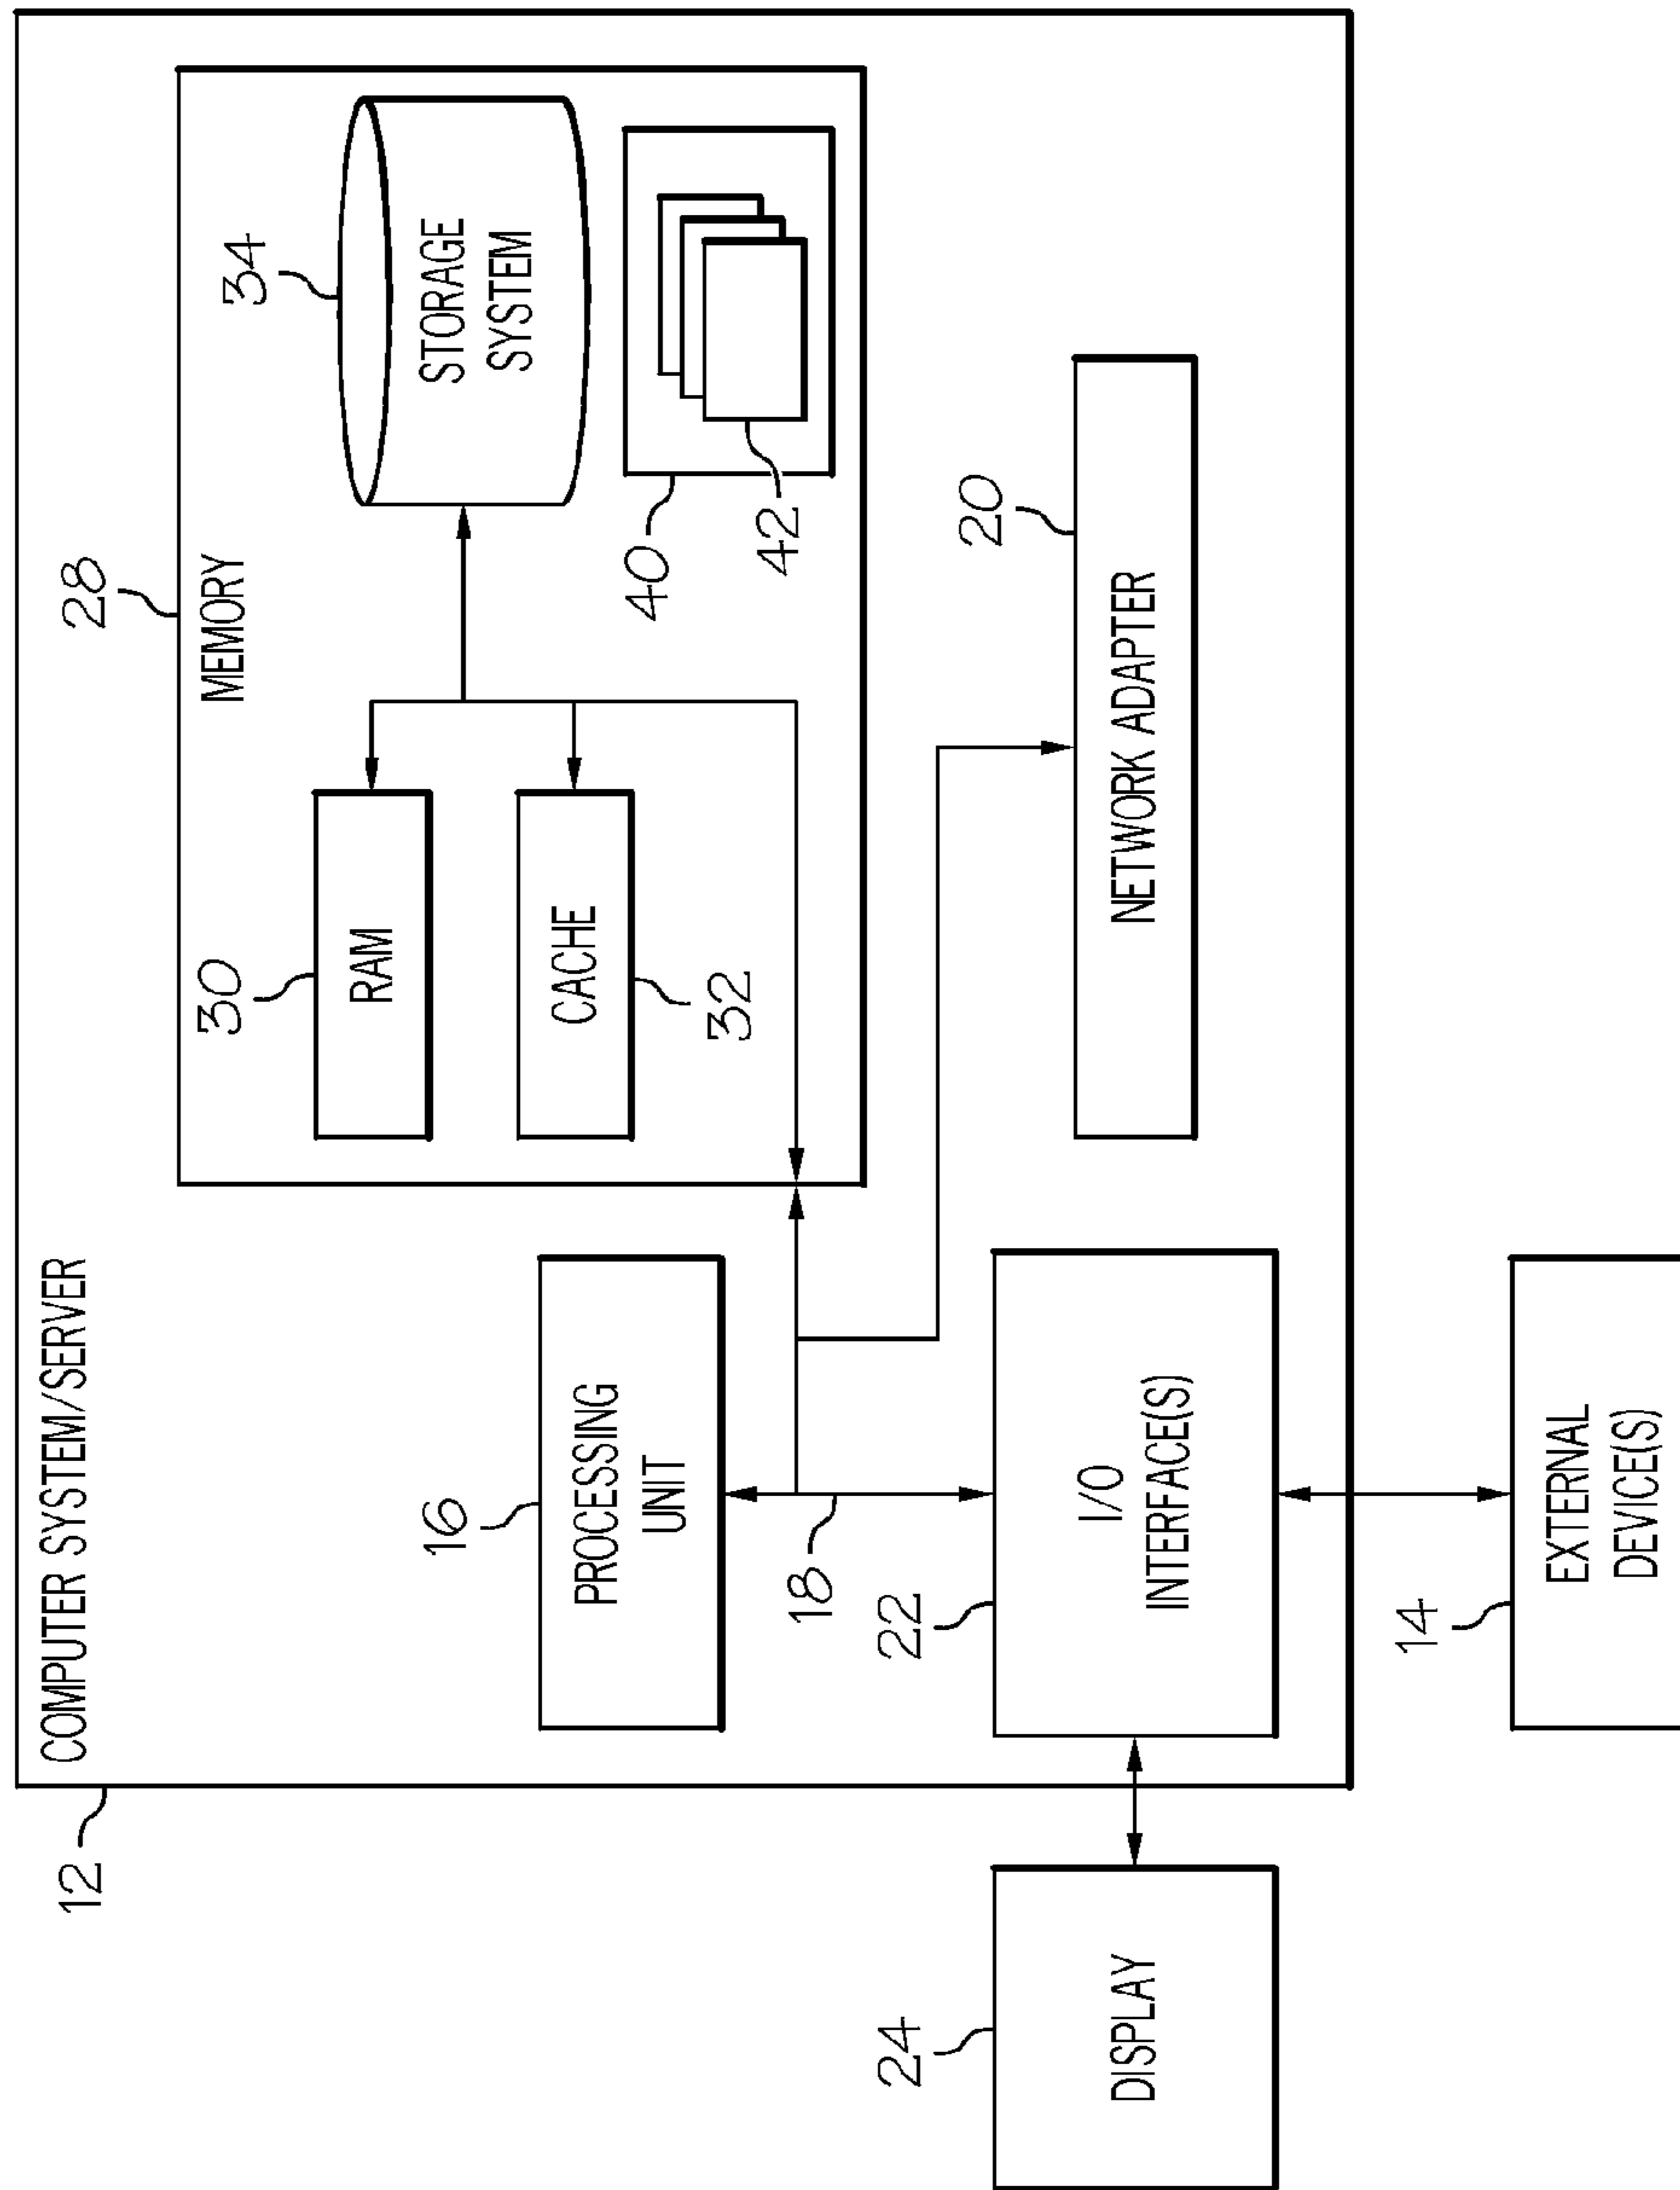


FIG. 1




FIG. 2

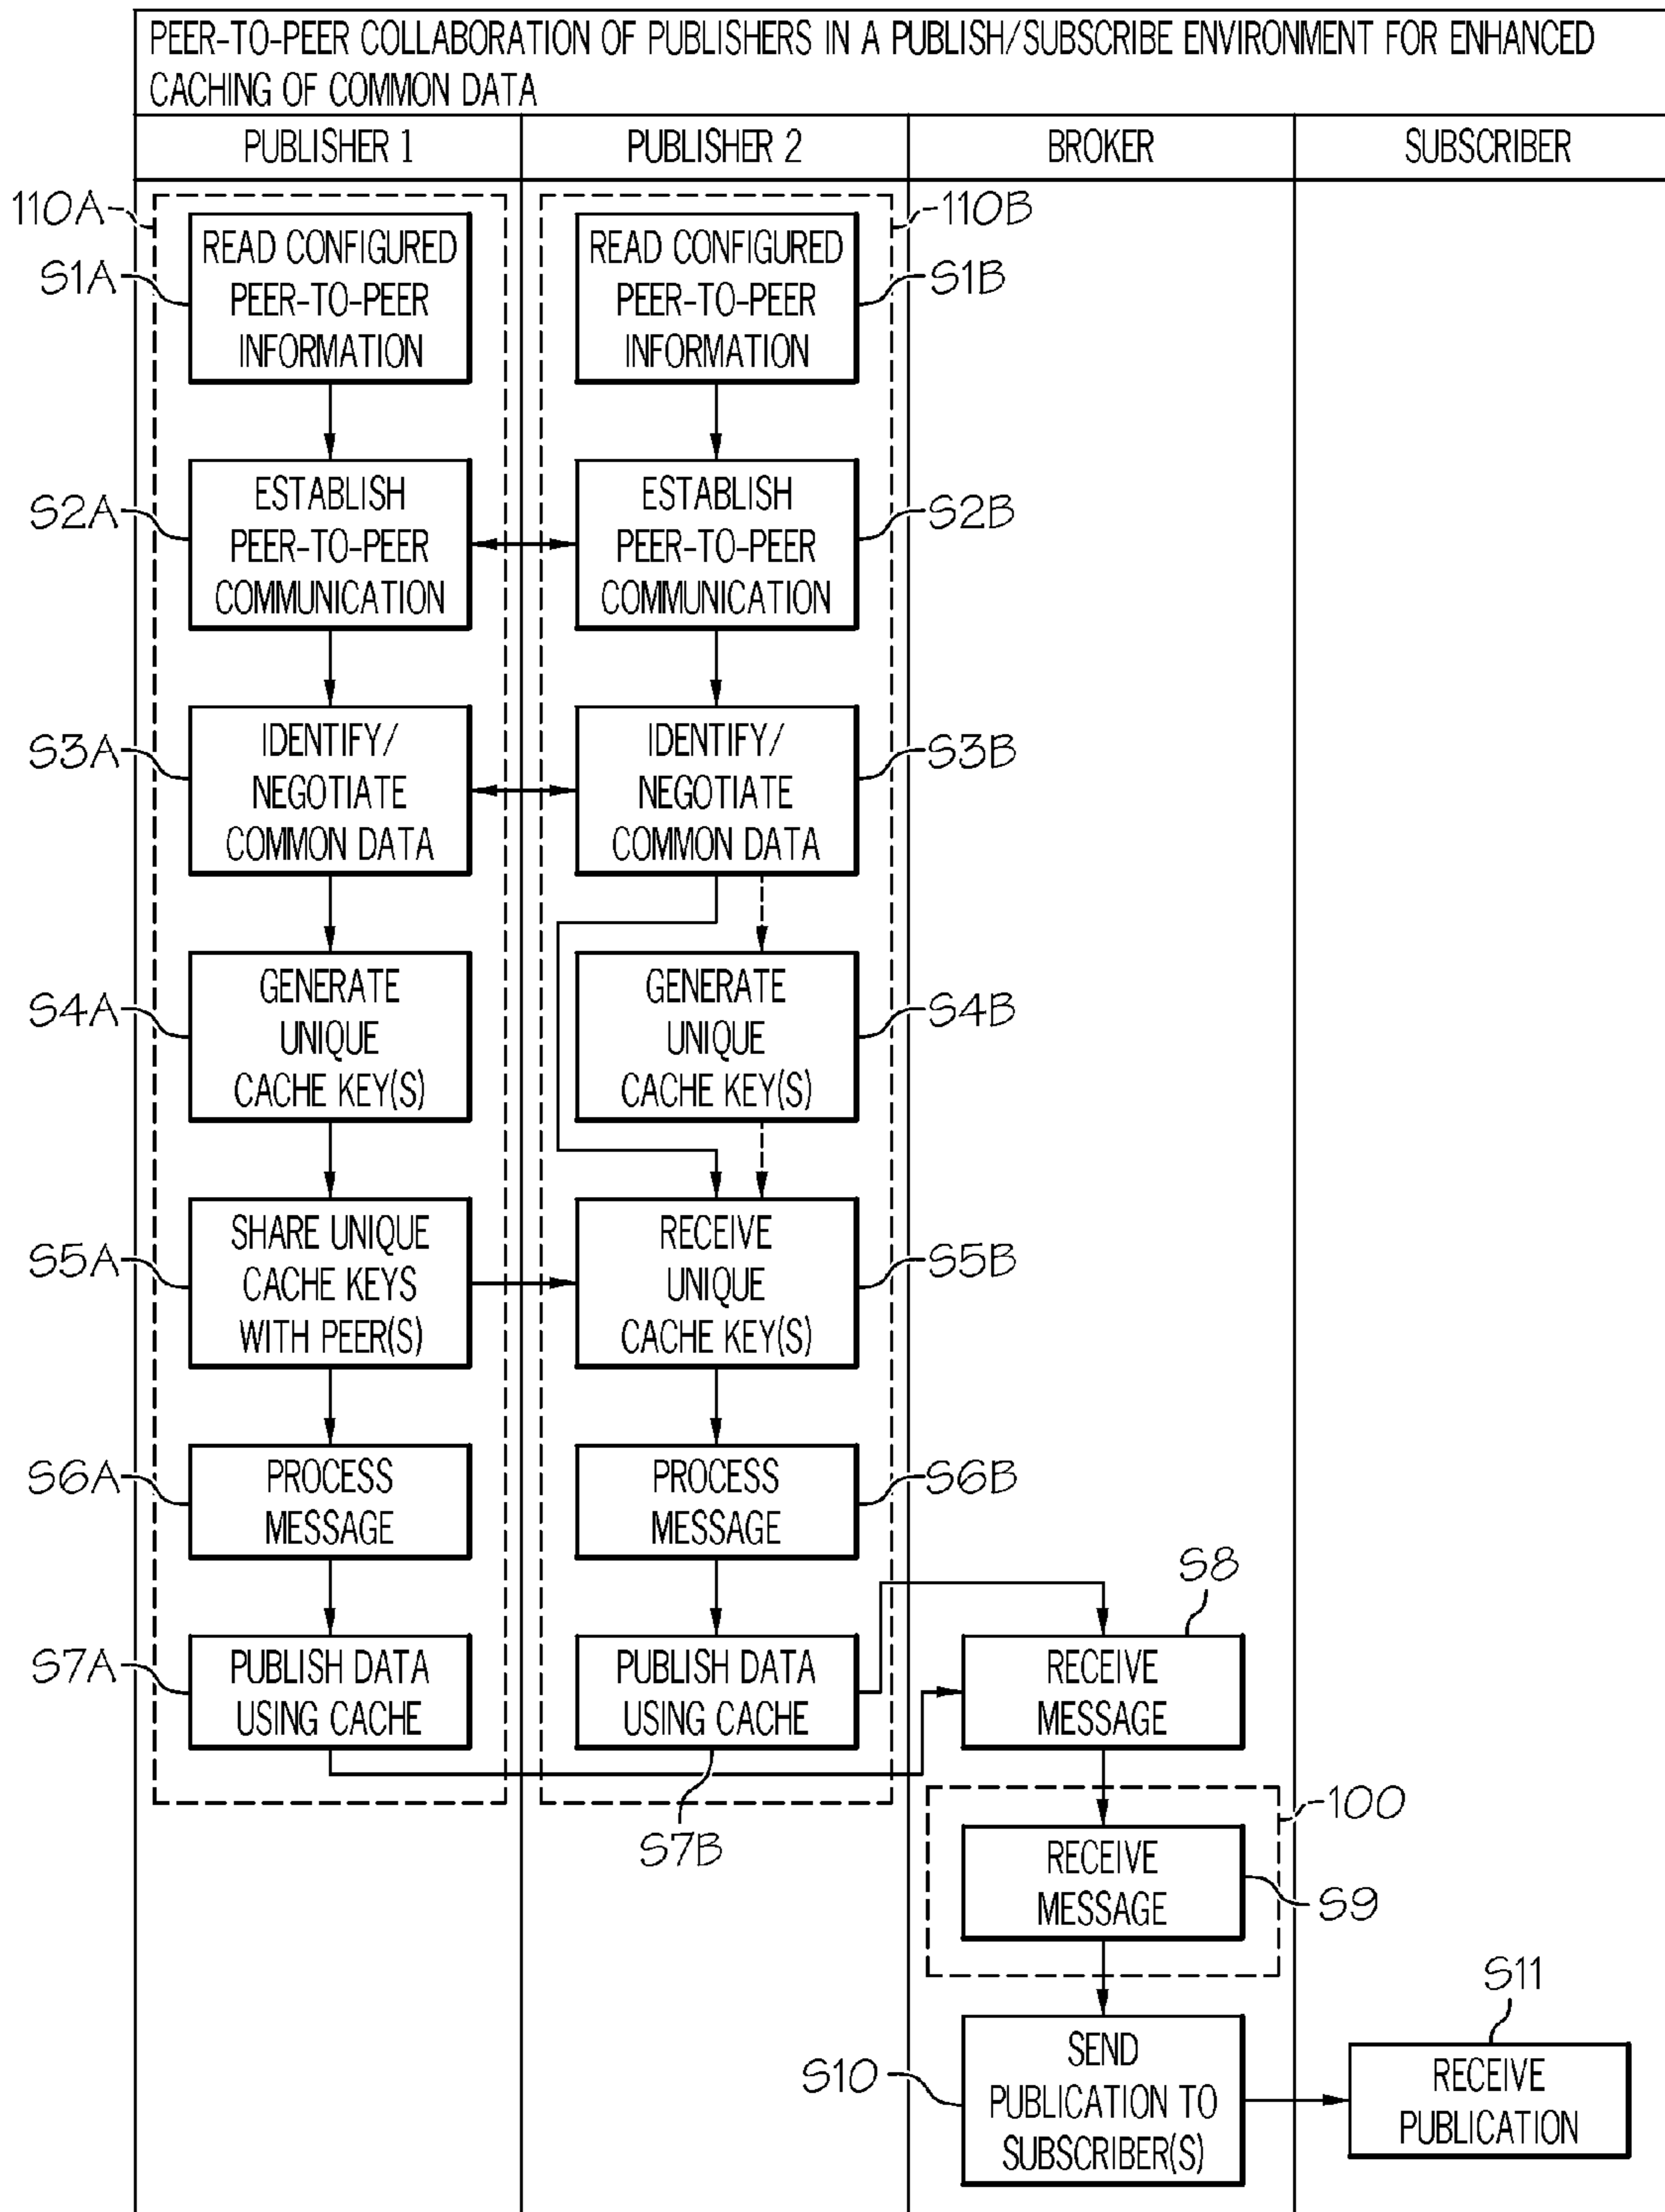


FIG. 3

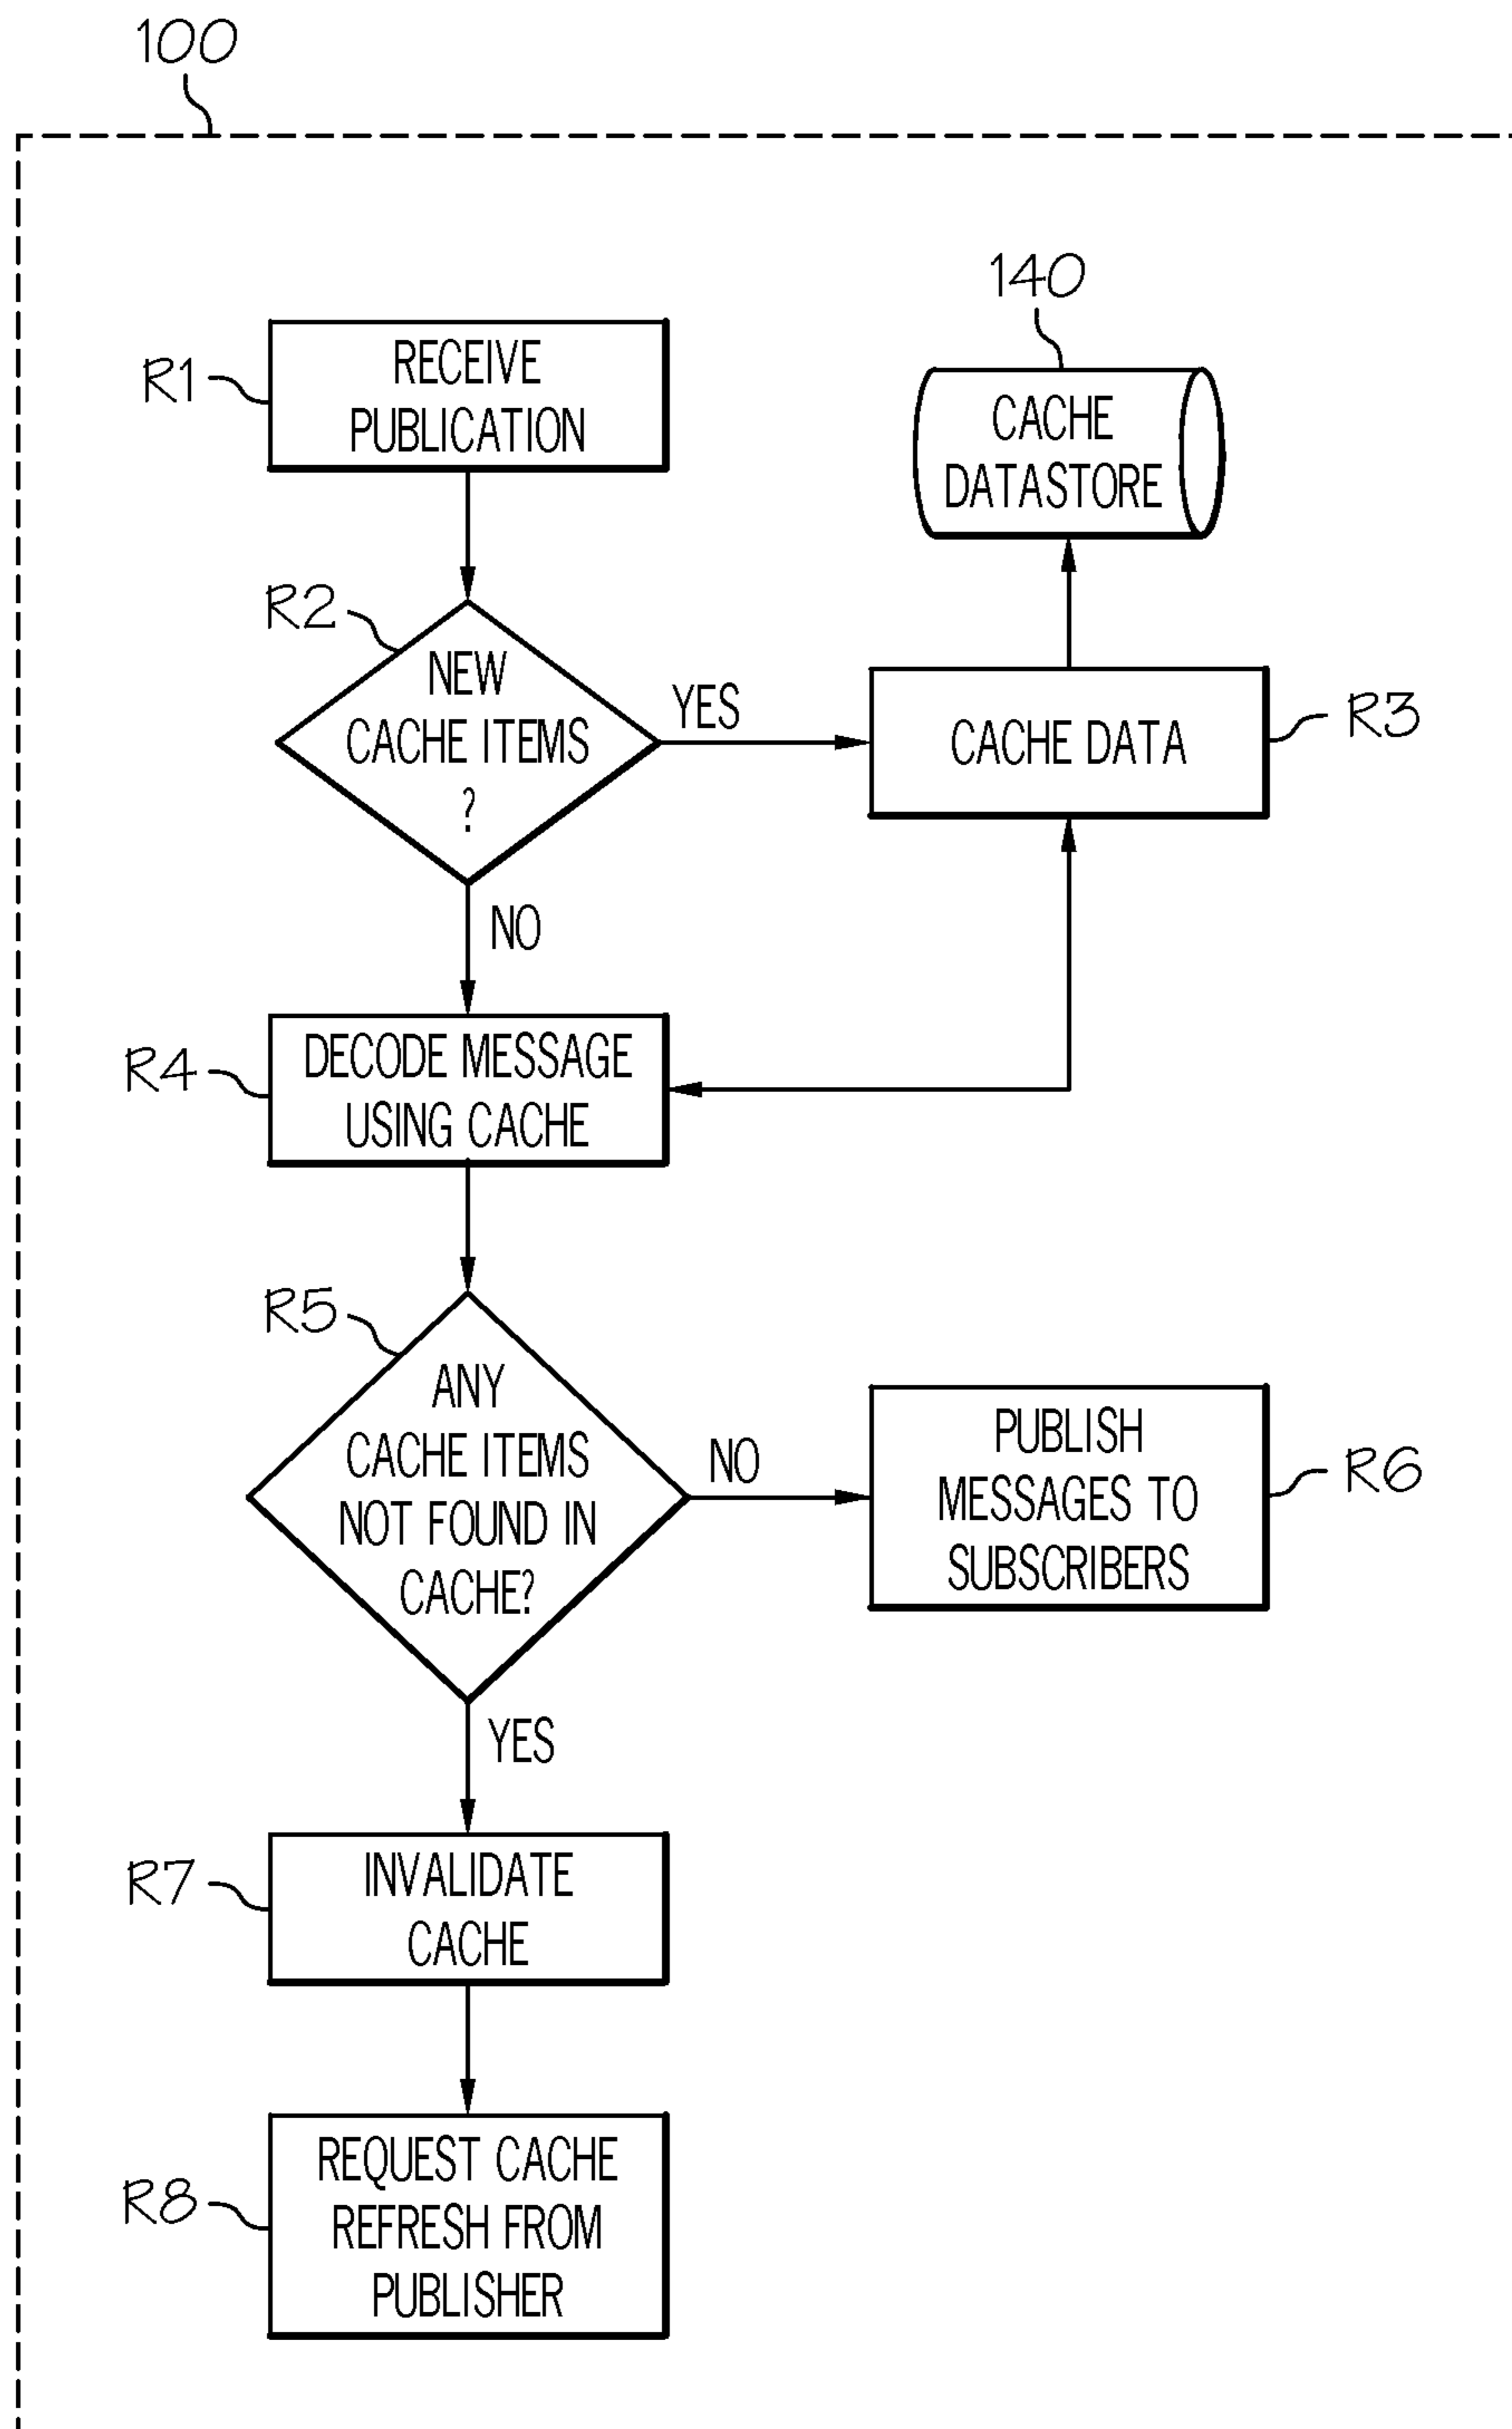


FIG. 4

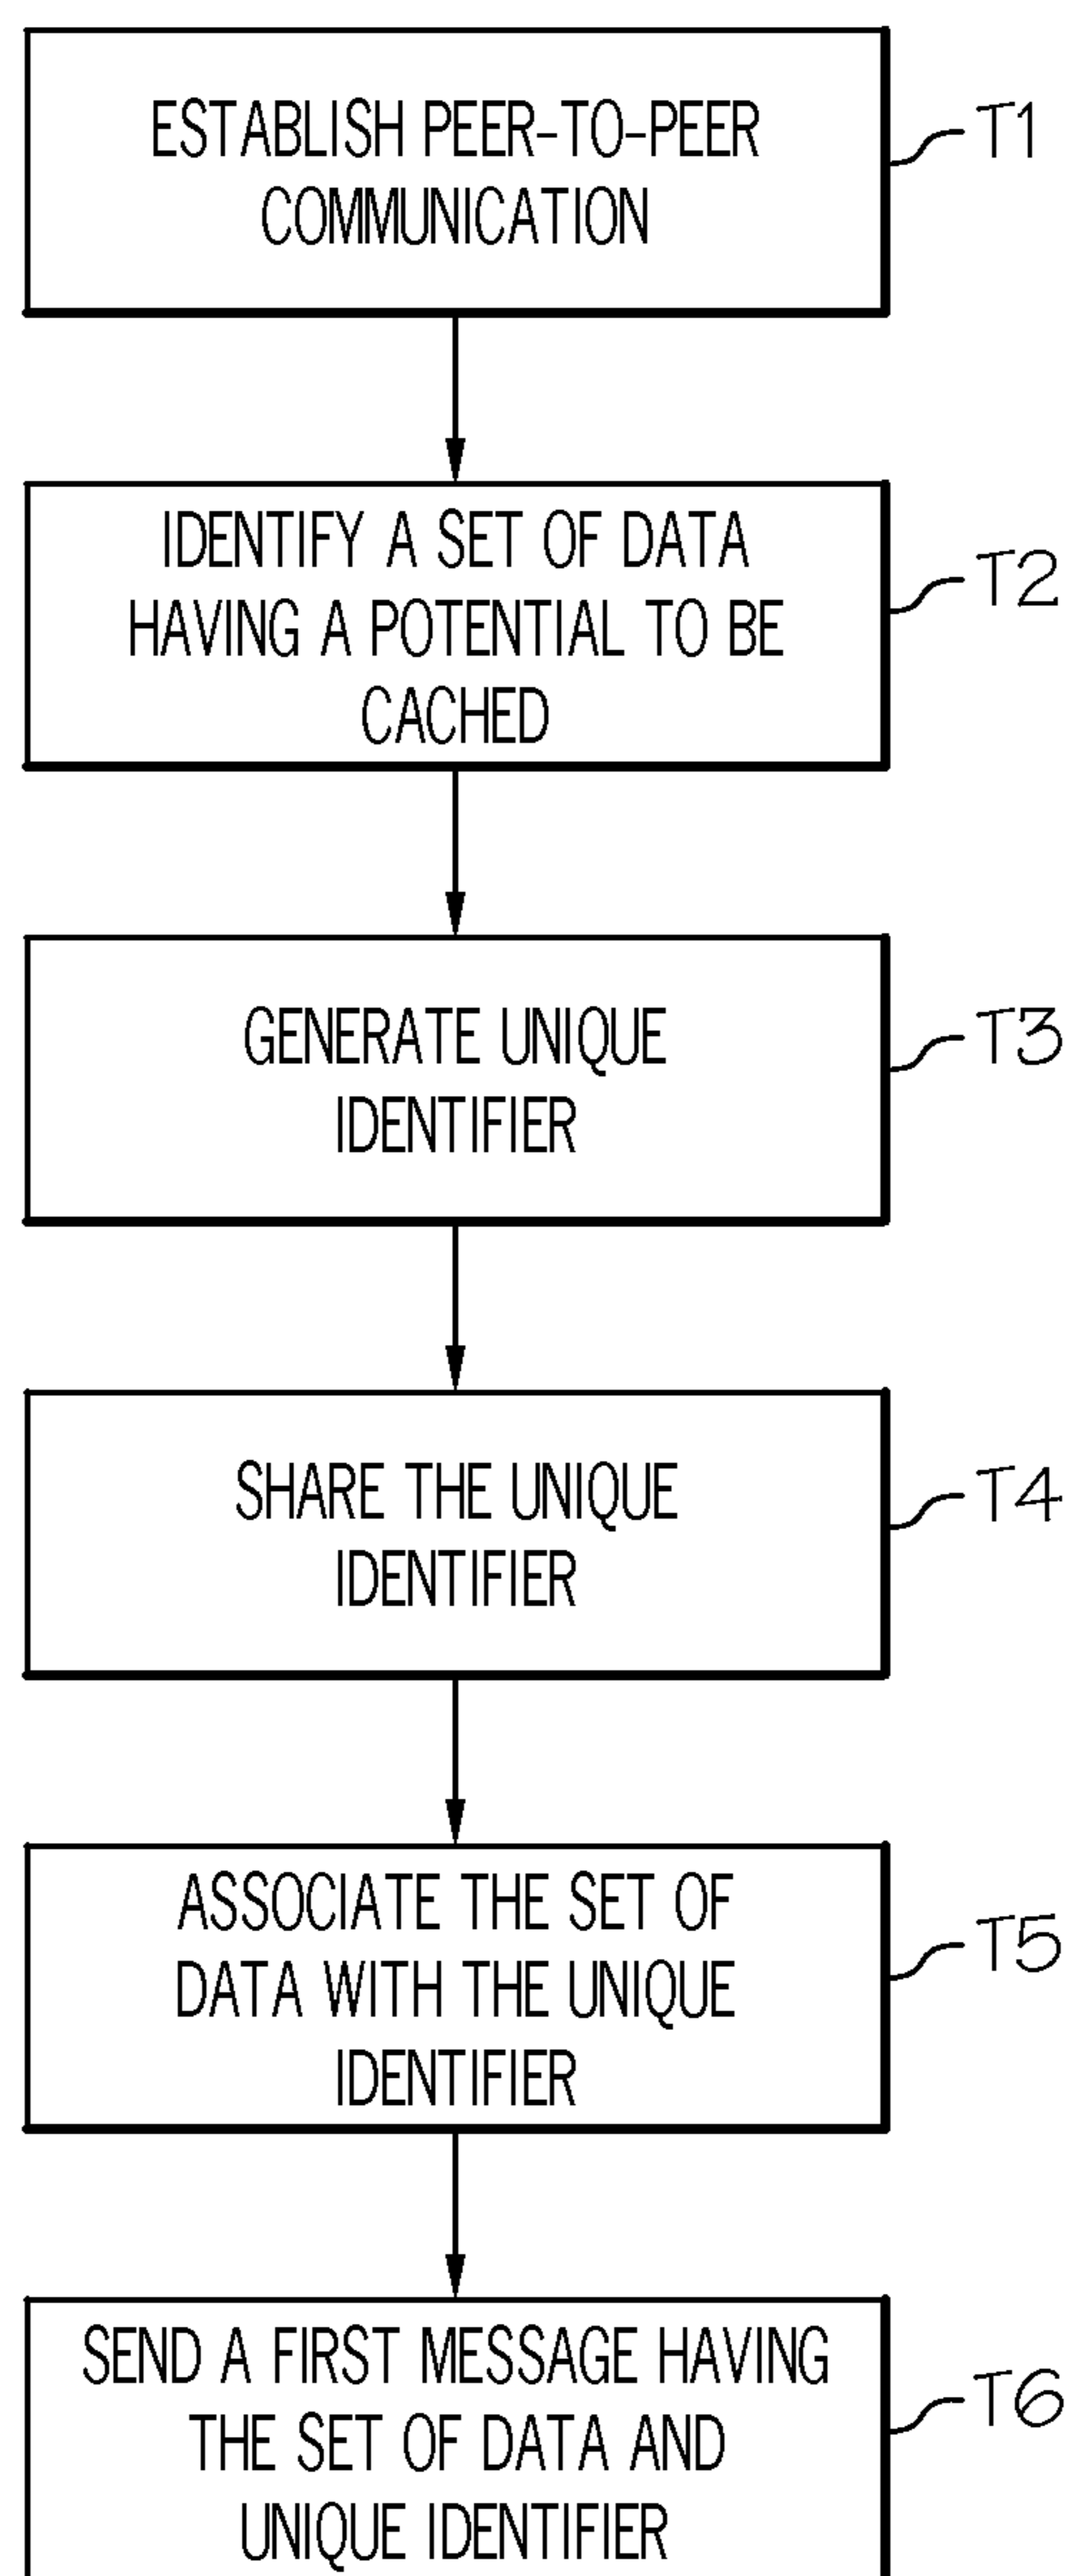


FIG. 5

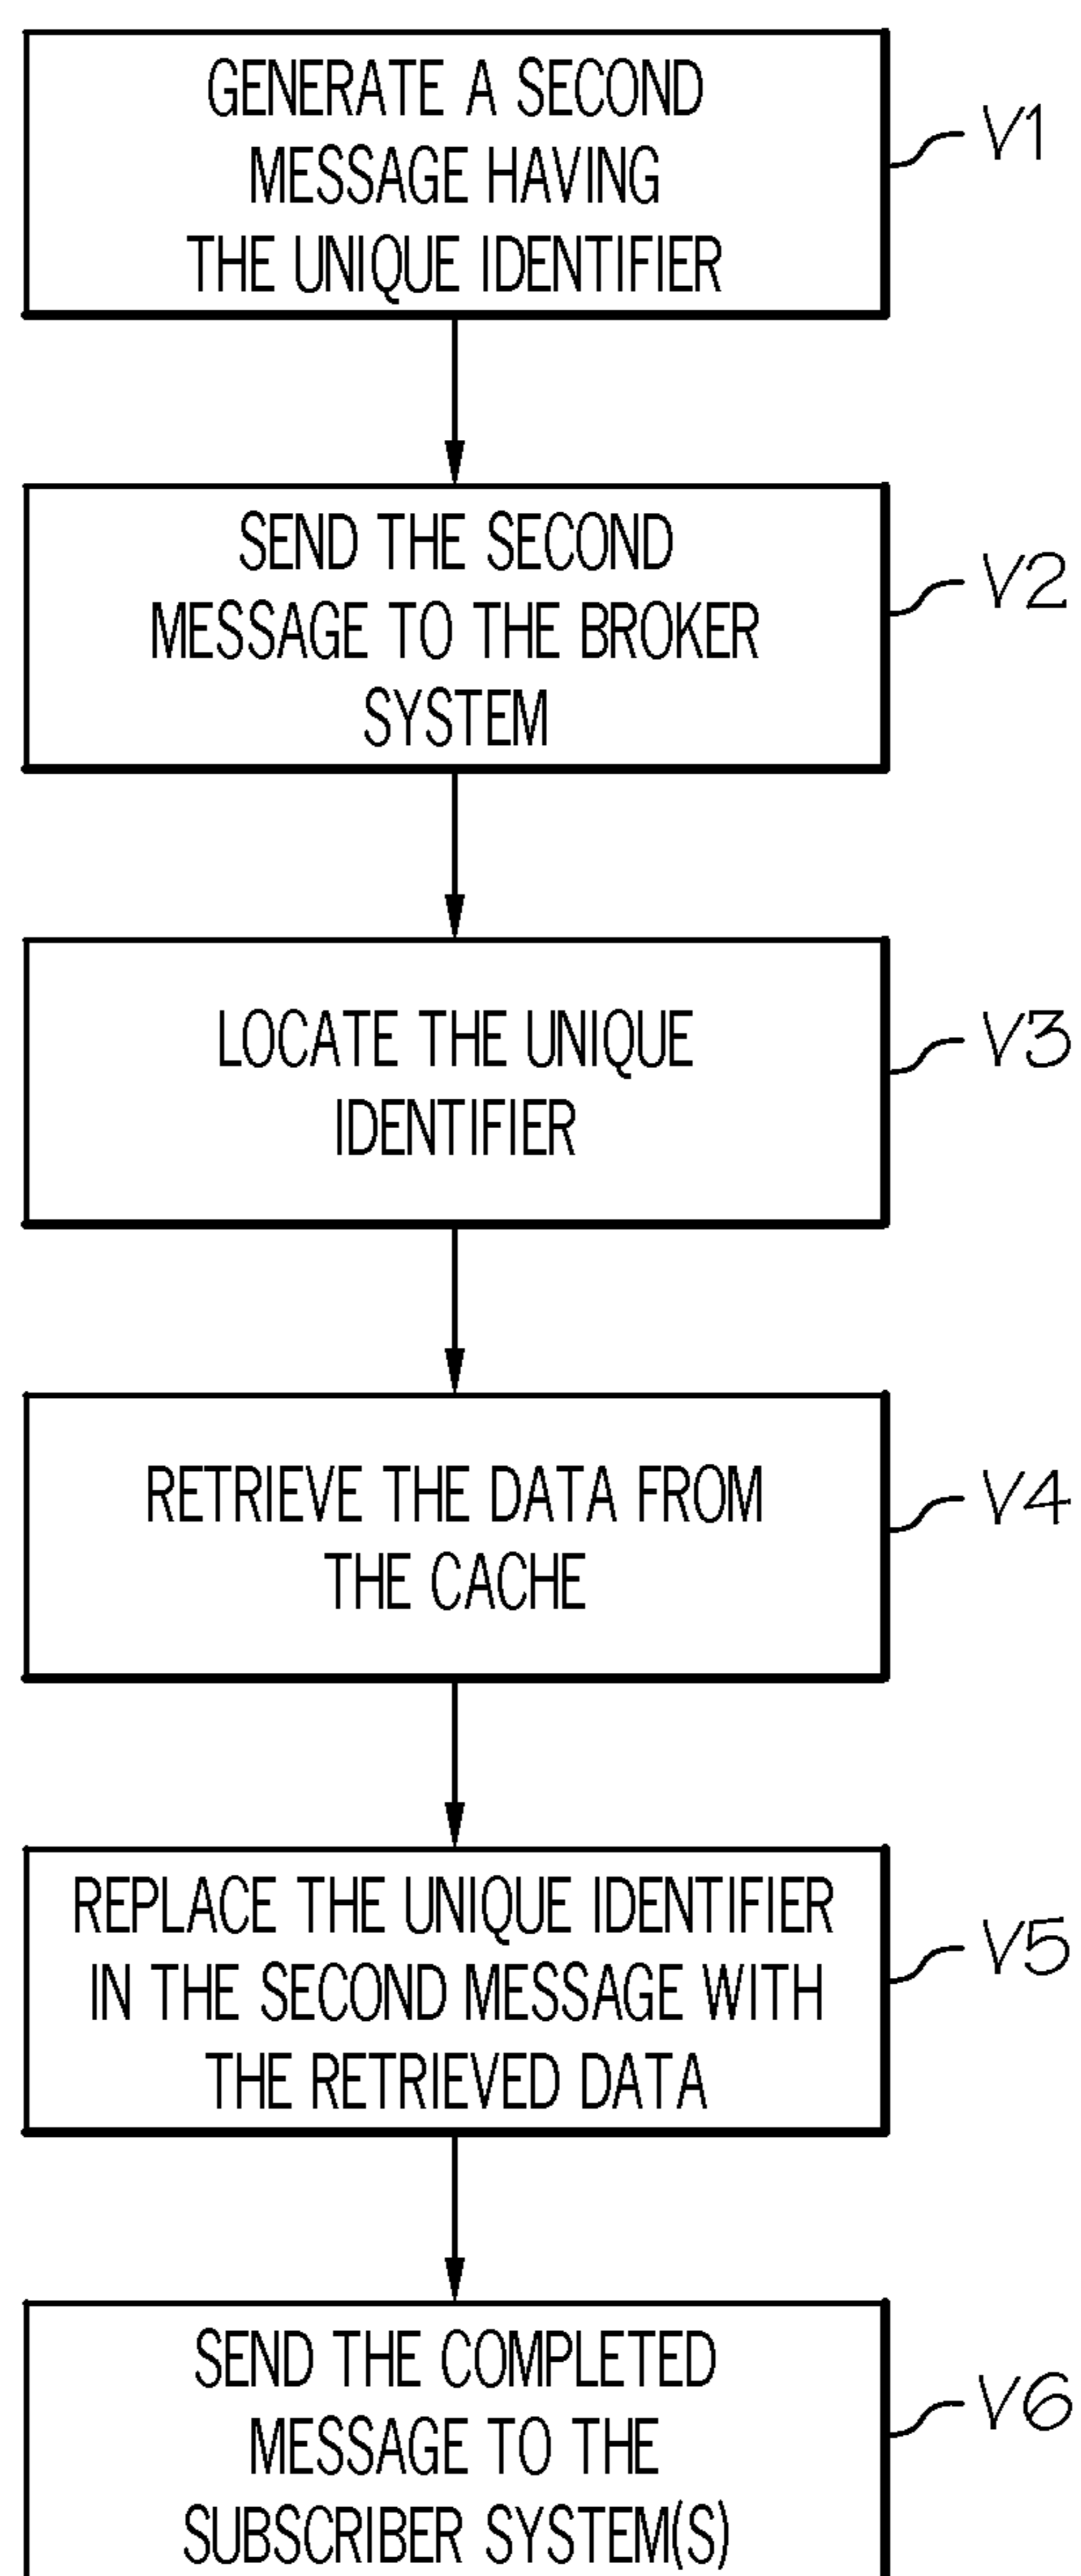


FIG. 6

**PEER-TO-PEER COLLABORATION OF
PUBLISHERS IN A PUBLISH-SUBSCRIPTION
ENVIRONMENT**

TECHNICAL FIELD

The present invention relates to publish-subscription (pub/sub) environments. Specifically, the present invention relates to an approach for caching potentially repetitive data in a pub/sub environment via peer-to-peer (P2P) collaboration of publishers.

BACKGROUND

Publish-subscription (pub/sub) messaging patterns are communication models used in many different industries today. For example, pub/sub is used in banking systems, weather systems, and other systems that require techniques to notify multiple entities of common information. In general, pub/sub is a messaging paradigm where senders (publishers) of messages typically are not programmed to send their messages to specific receivers (subscribers). Rather, published messages can be characterized into classes, without knowledge of what (if any) subscribers there may be. Subscribers express interest in one or more classes and can receive messages that are of interest, without knowledge of what publishers there are. This decoupling of publishers and subscribers can allow for greater scalability and a more dynamic network topology.

One deficiency in pub/sub models is that all of the content (i.e., the body/payload of the message) is sent to all users, and thus the amount of network traffic generated from the publisher is, in the simplest form, the size of the message multiplied by the number of subscribers. Moreover, it can be common for the same set of data to be re-sent multiple times (in different messages) to the same subscriber(s). As such, the volume of the data generated can be large and consume a great deal of network bandwidth.

SUMMARY

In general, embodiments of the present invention provide an approach for an enhanced pub/sub environment whereby a peer group of publishers is defined and can collaborate to determine when to cache data being sent to a common broker or brokers (e.g., identify data that is common to multiple publishers which may be cacheable). Specifically, in a typical embodiment, configuration data or the like will be read to define a peer group among the publisher systems in the pub/sub messaging environment. Once such a group has been identified, communication among the peers therein will be established. This communication will be utilized to identify a set of data that is common to the set (at least one) of publisher systems in the peer group, which has a potential to be cacheable. Once such data has been identified, a unique identifier/set of cache keys will be generated corresponding thereto, and associated with the set of data. A message comprising the set of data and corresponding unique identifier will then be sent to a broker system (e.g., from the publisher system(s)), which will store (or have stored) the set of data in a cache, and pass the underlying message along to applicable subscriber system(s). When a future message is to be generated with the set of data previously cached, the publisher system will instead substitute the unique identifier for the cached set of data to yield an abbreviated message, which is sent to the broker system. Upon receipt of the abbreviated message, the broker system will locate the unique identifier, retrieve the

corresponding set of data from cache, substitute the set of data for the unique identifier to yield a completed message, and send the completed message to the applicable subscriber system(s).

5 A first aspect of the present invention provides a method for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, comprising: establishing P2P communication between a set of publisher systems in the pub/sub messaging environment; 10 identifying a set of data that is common to the set of publisher systems having a potential to be cached; generating a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable; sharing the unique identifier among the set of publisher systems; associating the set of data with the unique identifier; and sending a first message having the set of data and the unique identifier to a broker system in the pub/sub messaging environment.

15 A second aspect of the present invention provides a system for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, comprising: at least one bus; at least one processor coupled to the at least one bus; and at least one memory 20 medium coupled to the at least one bus, the at least one memory medium comprising instructions to: establish P2P communication between a set of publisher systems in the pub/sub messaging environment; identify a set of data that is common to the set of publisher systems having a potential to 25 be cached; generate a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable; share the unique identifier among the set of publisher systems; associate the set of data with the unique identifier; and send a first message having the set of data and the unique 30 identifier to a broker system in the messaging pub/sub environment.

35 A third aspect of the present invention provides an at least one computer program product for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, the at least one computer program product comprising at least one computer readable storage media, and program instructions stored on the at least one computer readable storage media, to: establish P2P communication between a set of publisher systems in the pub/sub 40 messaging environment; identify a set of data that is common to the set of publisher systems having a potential to be cached; generate a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable; share the unique identifier among the set of publisher systems; associate the set of data with the unique identifier; and send a first message having the set of data and the unique identifier to a broker system in the pub/sub messaging environment.

45 A fourth aspect of the present invention provides a method for deploying a system for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, comprising: providing a computer infrastructure being operable to: establish P2P communication between a set of publisher systems in the pub/sub 50 messaging environment; identify a set of data that is common to the set of publisher systems having a potential to be cached; generate a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable; share the unique identifier among the set of publisher systems; associate the set of data with the unique identifier; and send a first message 55 60 65

having the set of data and the unique identifier to a broker system in the pub/sub messaging environment.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other features of this invention will be more readily understood from the following detailed description of the various aspects of the invention taken in conjunction with the accompanying drawings in which:

FIG. 1 depicts a computing node according to an embodiment of the present invention.

FIG. 2 depicts a pub/sub environment according to an embodiment of the present invention.

FIG. 3 depicts a component flow diagram according to an embodiment of the present invention.

FIG. 4 depicts another component flow diagram according to an embodiment of the present invention.

FIG. 5 depicts a method flow diagram according to an embodiment of the present invention.

FIG. 6 depicts another method flow diagram according to an embodiment of the present invention.

The drawings are not necessarily to scale. The drawings are merely schematic representations, not intended to portray specific parameters of the invention. The drawings are intended to depict only typical embodiments of the invention, and therefore should not be considered as limiting the scope of the invention. In the drawings, like numbering represents like elements.

DETAILED DESCRIPTION

Illustrative embodiments will now be described more fully herein with reference to the accompanying drawings, in which exemplary embodiments are shown. This disclosure may, however, be embodied in many different forms and should not be construed as limited to the exemplary embodiments set forth herein. Rather, these exemplary embodiments are provided so that this disclosure will be thorough and complete and will fully convey the scope of this disclosure to those skilled in the art. In the description, details of well-known features and techniques may be omitted to avoid unnecessarily obscuring the presented embodiments.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of this disclosure. As used herein, the singular forms “a”, “an”, and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. Furthermore, the use of the terms “a”, “an”, etc., do not denote a limitation of quantity, but rather denote the presence of at least one of the referenced items. It will be further understood that the terms “comprises” and/or “comprising”, or “includes” and/or “including”, when used in this specification, specify the presence of stated features, regions, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, regions, integers, steps, operations, elements, components, and/or groups thereof.

In general, embodiments of the present invention provide an approach for an enhanced pub/sub environment whereby a peer group of publishers is defined and can collaborate to determine when to cache data being sent to a common broker or brokers (e.g., identify data that is common to multiple publishers which may be cachable). Specifically, in a typical embodiment, configuration data or the like will be read to define a peer group among the publisher systems in the pub/sub messaging environment. Once such a group has been identified, communication among the peers therein will be

established. This communication will be utilized to identify a set of data that is common to the set (at least one) of publisher systems in the peer group, which has a potential to be cacheable. Once such data has been identified, a unique identifier/set of cache keys will be generated corresponding thereto, and associated with the set of data. A message comprising the set of data and corresponding unique identifier will then be sent to a broker system (e.g., from the publisher system(s)), which will store (or have stored) the set of data in a cache, and pass the underlying message along to applicable subscriber system(s). When a future message is to be generated with the set of data previously cached, the publisher system will instead substitute the unique identifier for the cached set of data to yield an abbreviated message, which is sent to the broker system. Upon receipt of the abbreviated message, the broker system will locate the unique identifier, retrieve the corresponding set of data from cache, substitute the set of data for the unique identifier to yield a completed message, and send the completed message to the applicable subscriber system(s).

The approaches described herein provide a solution to optimize network traffic between publishers and brokers and provide a system that works in complex production environments commonly found where only uni-directional communication is available between publishers and brokers. Similarly, in low latency networks, it may not be ideal to have two-way communications between the publisher and broker.

It is noted that caching of data hereunder could occur via the publisher system, the broker system and/or the subscriber system. Along these lines, the location of the cache is not intended to be limiting.

Referring now to FIG. 1, a schematic of an example of a computing node is shown. Computing node 10 is only one example of a suitable computing node and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention described herein. Regardless, computing node 10 is capable of being implemented and/or performing any of the functionality set forth hereinabove.

In computing node 10, there is a computer system/server 12, which is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system/server 12 include, but are not limited to, personal computer systems, server computer systems, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed computing environments that include any of the above systems or devices, and the like.

Computer system/server 12 may be described in the general context of computer system-executable instructions, such as program modules, being executed by a computer system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server 12 may be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

As shown in FIG. 1, computer system/server 12 in computing node 10 is shown in the form of a general-purpose

computing device. The components of computer system/server **12** may include, but are not limited to, one or more processors or processing units **16**, a system memory **28**, and a bus **18** that couples various system components including system memory **28** to processor **16**.

Bus **18** represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnects (PCI) bus.

Computer system/server **12** typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server **12**, and it includes both volatile and non-volatile media, removable and non-removable media.

System memory **28** can include computer system readable media in the form of volatile memory, such as random access memory (RAM) **30** and/or cache memory **32**. Computer system/server **12** may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system **34** can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM, or other optical media can be provided. In such instances, each can be connected to bus **18** by one or more data media interfaces. As will be further depicted and described below, memory **28** may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

The embodiments of the invention may be implemented as a computer readable signal medium, which may include a propagated data signal with computer readable program code embodied therein (e.g., in baseband or as part of a carrier wave). Such a propagated signal may take any of a variety of forms including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium including, but not limited to, wireless, wireline, optical fiber cable, radio-frequency (RF), etc., or any suitable combination of the foregoing.

Program/utility **40**, having a set (at least one) of program modules **42**, may be stored in memory **28** by way of example, and not limitation, as well as an operating system, one or more application programs, other program modules, and program data. In general, program/utility **40** performs the function of the present invention as described herein. Each of the operating systems, one or more application programs, other program modules, and program data or some combination thereof, may include an implementation of a networking environment. Program modules **42** generally carry out the functions and/or methodologies of embodiments of the invention as described herein. Computer system/server **12** may also

communicate with one or more external devices **14** such as a keyboard, a pointing device, a display **24**, etc.; one or more devices that enable a consumer to interact with computer system/server **12**; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server **12** to communicate with one or more other computing devices. Such communication can occur via I/O interfaces **22**.

Still yet, computer system/server **12** can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter **20**. As depicted, network adapter **20** communicates with the other components of computer system/server **12** via bus **18**. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server **12**. Examples include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

Embodiments of the present invention provide an approach for a publisher-based cache that can be utilized to reduce the volume of data (network traffic) delivered to individual subscribers in a pub/sub system.

Referring now to FIG. **2**, a pub/sub environment **70** in accordance with at least one embodiment of the present invention is shown in greater detail. As depicted, environment **70** comprises a publisher system **72** (also referred to below as "publisher") in communication with a subscriber system **74** (also referred to below as "subscriber") via a broker system **76** (also referred to below as "broker"). Systems **72**, **74**, and **76** communicate via network links **78A-N**, which can represent any type of network links and/or communication networks now known or later developed. It is understood that publisher system **72**, subscriber system **74**, and broker system **76** each comprise one or more computer systems **12A-C** (e.g., which comprises components/elements similar to computer system **12** of FIG. **1**). As such, publisher system **72**, subscriber system **74**, and broker system **76** can each comprise a program/utility **40A-C** that enable corresponding functions thereof. Specifically, publisher system **72** comprises publisher program **40A** having a set (at least one) of publisher program modules **42A**, broker system **76** comprises broker program **40B** having a set of broker program modules **42B**, and subscriber system **74** comprises a subscriber program **40C** having a set of subscriber program modules **42C**. Program modules **42A-C** of programs **40A-C** typically enable/perform the functions that will be described below in conjunction with FIGS. **3-5**.

Referring now to FIG. **3**, a component flow diagram illustrating the process involved with a typical embodiment of the present invention is shown in greater detail. It is understood in advance that where identical steps are shown as being performed by multiple publishers, only a single instance may need to be performed. Identical steps are shown in parallel to illustrate that the cache engines **110A-B** of either publisher system could carry out the underlying functionality. In steps **S1A-B**, publisher systems first learn what peers are in their peer group. This may be done via a configuration file that is read or some other means by which the publisher learns about its peers. In one embodiment, an administrator or the like could provide this information when the publisher is created. The publisher systems in a peer group will establish communication between its peer(s) in steps **S2A-B**. This may be a simple one to one relationship, one-to-many, or many-to-many. The peer group may define a master peer or some other hierarchy for managing communications between the peer group. As a publisher system runs, either by itself or in con-

junction with the peer group, the publisher system will determine what data is publishing, which should be cached in steps S3A-B.

Along these lines, the identification process can be configured with attributes such as one or more of the following:

If data is going to be in more than "N" number of messages;

If data is going to be in more than "M" number of messages and is over X" number of bytes.

These determinations can be made based upon historical information of messages already published or about the types of data typically published (e.g., common extensible markup language (XML) schemas that are known to be published may have common, lengthy tags or sequences of tags that can be cached).

The collaboration between the peer group is then used to further determine what data should be cached. For instance, if the rule for an individual publisher states that data should be cached if the data has occurred in ten or more past messages, the peer group may define a rule that states if the data has occurred at least twice, historically check with the peer group and see if, as a group, the data is seen ten or more times (or some other configurable value). If those criteria have been met, then have the peer group cache this data. Likewise, the default rule may be that if an individual publisher system determines that specific data should be cached that all publishers in the peer group should be notified and always use the cache keys/unique identifiers associated with the data. Along these lines, the cache keys/unique identifier is established for the data in step S4A or S4B. In this step, it is not necessary for more than one publisher system to generate the unique identifier. As such, only step S4A or S4B need to be performed. Along these lines, the generation of cache keys/unique identifiers could be done by a master peer in the group, negotiated between peers, or some other method. Regardless, once the cache key/unique identifier is created and associated with the data, it is shared among all peers in the peer group in steps S5A-B.

In steps S6A-B the message is processed. This typically involves tagging/annotating/associating the cacheable data in the message with the cache keys/unique identifier, and need only be performed by a single publisher system. This process could be implemented using one of the methods described in the sections below entitled: Potential Implementations I and II. Once the message has been processed, it is published in step S7A or S7B, and received by the broker system in step S8. Upon receipt, the broker system will process the message in step S9, and send the message to applicable subscriber system(s) in step S10, which will be received in step S11.

The processing of messages in step S9 by the broker system cache engine 100 will be described in greater detail with respect to FIG. 4. However, in general, the broker system cache engine 100 will detect and locate the unique identifier within the message, and store (or have stored) the corresponding data in a cache (e.g., a broker system cache, a subscriber system cache, etc.). Then, when a message is later received having the unique identifier (e.g., an abbreviated message is received), the broker system cache engine 100 will retrieve the corresponding data, replace the unique identifier in the message with the retrieved data to yield a completed message, and then send the completed message to the applicable subscriber system(s).

Although not shown, the subscriber can have access to the broker system cache and/or maintain its own cache. This allows the subscriber to maintain or have access to previously cached data. Regardless, when the publisher is creating a message with previously cached data, the publisher system will substitute the unique identifier for such data to yield an

abbreviated message, and then send the abbreviated message to the broker system. Upon receipt, the broker system can then locate the unique identifier in the abbreviated message, retrieve the associated data from the cache, replace the unique identifier in the message with the retrieved data, and then send the completed message to the applicable subscriber system(s). Alternatively, in the event the subscriber system(s) maintain or have access to the cached data, the abbreviated message received by the broker system from the publisher system can be sent to the subscriber system(s), which can retrieve the previously cached data that can be used by the subscriber system(s) to complete the message (e.g., via a subscriber cache engine).

Referring now to FIG. 4, the functions of broker cache engine 100 will be described in greater detail. In step R1, a message/publication is received from the publisher system. In step R2, it is determined whether new items need to be cached. If so, they are stored in data store/cache 140 in step R3. Once all cacheable data items have been so stored, the message is analyzed to see if it includes unique identifier(s) corresponding to data cached from previous messages (e.g., the message is an abbreviated message) in step R4. If so, that data is retrieved from cache 140 and inserted into the message. Once all new cacheable data has been cached, and previously cached data has been retrieved, it is determined in step R5 whether or not the message included cache keys/unique identifiers (pertaining to previously cached data) that were not valid (e.g., the cache keys/unique identifiers did not appear to be associated with any data stored in cache 140). If so, the cache can be invalidated in step R7, and a refresh of cache 140 can be requested from the publisher system in step R8. Regardless, once any broker-side processing is complete, the message can be sent to the subscriber system(s) in step R6. As indicated above, the subscriber system(s) can receive either a completed message (e.g., based on step R4) or an abbreviated message in the event the subscriber system(s) have access to previously cached data as indicated above.

It is understood that it may be beneficial to allow a broker system to invalidate cache items on a publisher system and/or a subscriber system. If, for example, the broker's cache engine can determine that data previously cached is never used again, or has a very low likelihood of being used again, it may send an invalidation message upon the next publish to the other system(s). This may be implemented as a custom header that contains the key and an invalidate directive. The subscriber system and/or publisher system, upon receiving and reading the message, will simply remove the cached item from its cache.

Additional aspects of the system:

The cache 140 (FIG. 4) illustrated above can be either in-memory or an on-disk data store (e.g. a database such as an IBM® DB2® database, flat file, binary file, binary encrypted file, etc).

The broker system's cache engine 100 may attach an expiration for each key. This value could be used by the subscriber system's or publisher system's caches to invalidate cache entries. This may be useful to limit the growth of the cache.

The broker system could send a subsequent message (to both publisher systems and subscriber systems) that would invalidate certain cache entries (e.g., it determines that data will never be used again).

The broker system's cache engine 100 can support automatic cache creation as described earlier. For example, if the last "N" number of messages on a topic contains a common data segment, then that data segment should be

marked for caching and a cache key/data pair created and sent with the next message.

Illustrative Implementations

Illustrative Implementation I

The cache engine either deletes or substitutes the cached data with the appropriate cache key/unique identifier (the system that allows for the data to be deleted is described below). For example, assume that the message being sent appears as follows:

```
AAAAAABBBBBBBBCCCCCDDDDDDDEEEEEEE
```

The cache engine states that the “B’s” will likely be used in future messages and instructs recipients that they should retain that portion in their local cache. In this example, the portion of data BBBB BBB will be referred to as ~rp1. The recipient will store the portion BBBB BBB and analyze incoming messages for the presence of ~rp1. Further assume that at some future point in time, a publisher wants to send a message that appears as follows:

```
FFFFFFFFBBBBBBBUUUUUJJJJJJ
```

Rather than sending this message, it can now send the following:

```
FFFFFFFF~rp1 UUUUUJJJJJJ
```

This allows for the sender to send less traffic on the network. Upon receipt of this message, the recipient will then reconstruct the message based on translation information from its cache and create the following message:

```
FFFFFFFFBBBBBBBUUUUUJJJJJJ
```

after expanding ~rp1 to BBBB BBB.

Illustrative Implementation II

If it was desired for cached data to be inserted at the 12th position of the message, for example, the message and a property/attribute (or custom header) would be generated that indicates rp1=12,x,y (where rp1 is the cache key/unique identifier, and 12 is the first position to insert the data, x is the next, y is the next, and so on). This method obviates the need to delete any un-cached message data that matches a cache indicator.

It is noted that multiple approaches may be included in these solutions.

Automatic Selection of Cache Data: In a typical embodiment, the proposed invention automatically determines which information should be cached.

Several techniques may be used to identify cacheable information. The techniques identified below may be used independently or in conjunction with each other.

Differencing Engine: This system analyzes messages either within a single topic, or across multiple topics capturing the similarities between messages. The most similar aspects of the messages are presumed good candidates for future caching. The similarities captured may be persisted across restarts such that caching efficiency is not interrupted. The items cached may change over time as the message profiles change. For example, if one type of message is dominant for a while, then a second disparate message type is dominant. Over time, cache entries related to the first type may be invalidated to make room for the second type.

Rate-Based Selection: It is common for some messages and topics to have a higher rate of publishing than others. This rate may be related to the importance or the dynamic nature of the subject for each particular topic. Topics with a higher publication rate should be selected for caching before topics with a lower publication rate. Rate-based selection counts the total number of messages per topic, or message types within a single topic, and selects the most published topics or message types for message portion caching.

Message Correlation Selection: Such a system may determine correlations between related topics and therefore cache

the potential similarities for the related topics upon detection of the first message. For example, there may be related topics of/scores/[court]/and/stats/[court] where court is a changing topic related to particular players on a particular cart. Every message published to a/stats/[court] is preceded by a message published to/scores/[court]. Both sets of messages may identify the player names for that particular court. Therefore, this method may discern that for every message published to/scores/[court] and/stats/[court], while components of the message change (current score, current stats), the player names stay the same. Therefore, for a short period of time after a message is published to/scores/[court] the player name component is eligible for caching until the corresponding/stats/[court] message is sent.

Finally, it is noted that different methods of signifying data entities may be used by both senders and recipients of data. These include universal hashing, hashing with checksums, rolling hashes, or even cryptographic hash functions, should security dictate. Each of these approaches offers certain advantages and can be used to speed and ease the implementation of any concepts described herein.

Referring now to FIG. 5, a method flow diagram according to an embodiment of the present invention is shown. In step T1, P2P communication is established between a set of publisher systems in the pub/sub messaging environment. In step T2, a set of data that is common to the set of publisher systems having a potential to be cached is identified. In step T3, a unique identifier corresponding to the set of data is generated on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable. In step T4, the unique identifier is shared among the set of publisher systems. In step T5, the set of data is associated with the unique identifier. In step T6, a first message having the set of data and the unique identifier is (generated and) sent to the broker system in the pub/sub messaging environment.

Referring now to FIG. 6, the retrieval of previously cached data will be described in greater detail. In step V1, a second message (as opposed to the first message of step T6 of FIG. 5) is generated on the publisher system. In this instance, the second message will have the unique identifier in place of a set of data that was previously cached. In step V2, the second message is sent to the broker system as an abbreviated message. In step V3, the unique identifier is located in the second message. In step V4, the set of data is retrieved from the cache, and used to replace the unique identifier in the second message in step V5. Then, in step V6, the second message is sent from the broker system to a set of subscriber systems. Alternatively, the abbreviated message received by the broker system could be sent to the subscriber system(s) which the subscriber system(s) could then utilize to retrieve the data from cache to complete the message. This option would occur where the subscriber system(s) have access to previously cached data.

While shown and described herein as a data caching solution, it is understood that the invention further provides various alternative embodiments. For example, in one embodiment, the invention provides a computer-readable/useable medium that includes computer program code to enable a computer infrastructure to provide data caching functionality as discussed herein. To this extent, the computer-readable/useable medium includes program code that implements each of the various processes of the invention. It is understood that the terms computer-readable medium or computer-useable medium comprise one or more of any type of physical embodiment of the program code. In particular, the computer-readable/useable medium can comprise program code embodied on one or more portable storage articles of manu-

11

facture (e.g., a compact disc, a magnetic disk, a tape, etc.), on one or more data storage portions of a computing device, such as memory 28 (FIG. 1) and/or storage system 34 (FIG. 1) (e.g., a fixed disk, a read-only memory, a random access memory, a cache memory, etc.).

In another embodiment, the invention provides a method that performs the process of the invention on a subscription, advertising, and/or fee basis. That is, a service provider, such as a Solution Integrator, could offer to provide data caching functionality. In this case, the service provider can create, maintain, support, etc., a computer infrastructure, such as computer system 12 (FIG. 1) that performs the processes of the invention for one or more consumers. In return, the service provider can receive payment from the consumer(s) under a subscription and/or fee agreement and/or the service provider can receive payment from the sale of advertising content to one or more third parties.

In still another embodiment, the invention provides a computer-implemented method for providing data caching functionality. In this case, a computer infrastructure, such as computer system 12 (FIG. 1), can be provided and one or more systems for performing the processes of the invention can be obtained (e.g., created, purchased, used, modified, etc.) and deployed to the computer infrastructure. To this extent, the deployment of a system can comprise one or more of: (1) installing program code on a computing device, such as computer system 12 (FIG. 1), from a computer-readable medium; (2) adding one or more computing devices to the computer infrastructure; and (3) incorporating and/or modifying one or more existing systems of the computer infrastructure to enable the computer infrastructure to perform the processes of the invention.

As used herein, it is understood that the terms “program code” and “computer program code” are synonymous and mean any expression, in any language, code, or notation, of a set of instructions intended to cause a computing device having an information processing capability to perform a particular function either directly or after either or both of the following: (a) conversion to another language, code, or notation; and/or (b) reproduction in a different material form. To this extent, program code can be embodied as one or more of: an application/software program, component software/a library of functions, an operating system, a basic device system/driver for a particular computing device, and the like.

A data processing system suitable for storing and/or executing program code can be provided hereunder and can include at least one processor communicatively coupled, directly or indirectly, to memory elements through a system bus. The memory elements can include, but are not limited to, local memory employed during actual execution of the program code, bulk storage, and cache memories that provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution. Input/output and/or other external devices (including, but not limited to, keyboards, displays, pointing devices, etc.) can be coupled to the system either directly or through intervening device controllers.

Network adapters also may be coupled to the system to enable the data processing system to become coupled to other data processing systems, remote printers, storage devices, and/or the like, through any combination of intervening private or public networks. Illustrative network adapters include, but are not limited to, modems, cable modems, and Ethernet cards.

The foregoing description of various aspects of the invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the

12

invention to the precise form disclosed and, obviously, many modifications and variations are possible. Such modifications and variations that may be apparent to a person skilled in the art are intended to be included within the scope of the invention as defined by the accompanying claims.

What is claimed is:

1. A method for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, comprising:

establishing P2P communication between a set of publisher systems in the pub/sub messaging environment; identifying a set of data that is common to the set of publisher systems having a potential to be cached; generating a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable; sharing the unique identifier among the set of publisher systems; associating the set of data with the unique identifier; sending a first message having the set of data and the unique identifier to a broker system in the pub/sub messaging environment; receiving the first message on the broker system; locating the unique identifier; causing the set of data to be stored in a cache; sending the first message to a set of subscriber systems in the pub/sub messaging environment; receiving a second message on the broker system from any of the set of publisher systems, the second message comprising an abbreviated message having the unique identifier in place of the set of data; locating the unique identifier in the abbreviated message; retrieving the set of data from the cache; replacing the unique identifier with the set of data in the second message to yield a completed message; and sending the completed message to the set of subscriber systems.

2. The method of claim 1, the associating and the sending being performed by the at least one of the set of publisher systems.

3. The method of claim 1, the cache being a subscriber system cache.

4. The method of claim 1, the cache being a broker system cache.

5. The method of claim 1, further comprising reading configuration data to identify the set of publisher systems as a peer group in the pub/sub messaging environment.

6. The method of claim 1, the identifying being based on at least one of the following factors: a quantity of the at least one future message that will include the set of data, and a size of the at least one future message.

7. A system for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, comprising:

at least one bus; at least one processor coupled to the at least one bus; and at least one memory medium coupled to the at least one bus, the at least one memory medium comprising instructions to: establish P2P communication between a set of publisher systems in the pub/sub messaging environment; identify a set of data that is common to the set of publisher systems having a potential to be cached; generate a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable;

13

share the unique identifier among the set of publisher systems;
 associate the set of data with the unique identifier;
 send a first message having the set of data and the unique identifier to a broker system in the messaging pub/sub environment;
 receive the first message on the broker system;
 locate the unique identifier;
 cause the set of data to be stored in a cache;
 send the first message to a set of subscriber systems in the messaging pub/sub environment;
 receive a second message on the broker system from any of the set of publisher systems, the second message comprising an abbreviated message having the unique identifier in place of the set of data;
 locate the unique identifier in the abbreviated message;
 retrieve the set of data from the cache;
 replace the unique identifier with the set of data in the second message to yield a completed message; and
 send the completed message to the set of subscriber systems.

8. The system of claim 7, the cache being a subscriber system cache.

9. The system of claim 7, the cache being a broker system cache.

10. The system of claim 7, the at least one memory medium further comprising instructions to:
 receive an invalidation message from the publisher system;
 and
 remove a set of items from the cache based on the invalidation message.

11. The system of claim 7, the identifying being based on at least one of the following factors: a quantity of the at least one future message that will include the set of data, and a size of the at least one future message.

12. At least one computer program product for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, the at least one computer program product comprising at least one non-transitory computer readable storage media, and program instructions stored on the at least one non-transitory computer readable storage media, to:
 establish P2P communication between a set of publisher systems in the pub/sub messaging environment;
 identify a set of data that is common to the set of publisher systems having a potential to be cached;
 generate a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable;
 share the unique identifier among the set of publisher systems;
 associate the set of data with the unique identifier;
 send a first message having the set of data and the unique identifier to a broker system in the pub/sub messaging environment;
 receive the first message on the broker system;
 locate the unique identifier;
 cause the set of data to be stored in a cache;
 send the first message to a set of subscriber systems in the pub/sub messaging environment;

14

receive a second message on the broker system from any of the set of publisher systems, the second message comprising an abbreviated message having the unique identifier in place of the set of data;
 locate the unique identifier in the abbreviated message;
 retrieve the set of data from the cache;
 replace the unique identifier with the set of data in the second message to yield a completed message; and
 send the completed message to the set of subscriber systems.

13. The at least one computer program product of claim 12, the cache being a subscriber system cache.

14. The at least one computer program product of claim 12, the cache being a broker system cache.

15. The at least one computer program product of claim 12, further comprising program instructions stored on the at least one non-transitory computer readable storage media to:
 receive an invalidation message from the publisher system;
 and
 remove a set of items from the cache based on the invalidation message.

16. The at least one computer program product of claim 12, the identifying being based on at least one of the following factors: a quantity of the at least one future message that will include the set of data, and a size of the at least one future message.

17. A method for deploying a system for establishing peer-to-peer (P2P) collaboration of publishers in a publish-subscription (pub/sub) messaging environment, comprising:
 providing a computer infrastructure being operable to:
 establish P2P communication between a set of publisher systems in the pub/sub messaging environment;
 identify a set of data that is common to the set of publisher systems having a potential to be cached;
 generate a unique identifier corresponding to the set of data on at least one of the set of publisher systems, the unique identifier identifying the set of data as cacheable;
 share the unique identifier among the set of publisher systems;
 associate the set of data with the unique identifier;
 send a first message having the set of data and the unique identifier to a broker system in the pub/sub messaging environment;
 receive the first message on the broker system;
 locate the unique identifier;
 cause the set of data to be stored in a cache;
 send the first message to a set of subscriber systems in the pub/sub messaging environment;
 receive a second message on the broker system from any of the set of publisher systems, the second message comprising an abbreviated message having the unique identifier in place of the set of data;
 locate the unique identifier in the abbreviated message;
 retrieve the set of data from the cache;
 replace the unique identifier with the set of data in the second message to yield a completed message; and
 send the completed message to the set of subscriber systems.

* * * * *